**PUBLIC**

SAP Adaptive Server Enterprise on AWS 16.0 SP04 PL01

Document Version: 1.0 – 2021-11-05

# Migrating Data to SAP ASE on AWS



THE BEST RUN SAP

# Content

# 1 Methods for Migrating Data from On-Premises to AWS

There are a number of ways to migrate your data from your on-premises data center to SAP ASE running on Amazon Web Services (AWS).

- `bcp` – copy data out of the on-premises data center and back into the SAP ASE on AWS.
- `dump` – dump the on-premises data, then upload the data to SAP ASE on AWS using the `load` command.
- `sybmigrate` utility.
- AWS Data Migration Service (DMS).

To migrate data, you need a licensed SAP ASE running on-premises and on AWS.

Choose a migration method based on your circumstances, such as the type of data you are migrating, the amount of data, and the size of your network connection.

| Method | Use This Method When: | Do Not Use This Method When: |
|---|---|---|
| `bcp` | • You are moving tables between the SAP ASE servers, or between an SAP ASE server and other data sources.<br>• The source and target database have different character sets.<br>• The source and target databases are on different platforms.<br><br>**i Note**<br>The source and target database versions need not be the same to use `bcp`. | • You need to create the schema or metadata along with data transfer.<br>• The data transfer must include database objects. |
| `dump` and `load` | • The device allocations on the target SAP ASE have equal or greater size than those of the source SAP ASE.<br>• The source and target databases have the same page size. | • The migration must load transaction dumps on different platforms.<br>• The dump from the source server uses a sort order that differs from the target server. |

| Method | Use This Method When: | Do Not Use This Method When: |
| --- | --- | --- |
| `sybmigrate` | You need to migrate the database objects along with the data.<br><br>**Note**<br>The source and target database platforms, database versions, and page size need not be the same to use `sybmigrate`. | There is any activity on the source or target server during the migration. If the objects are created, modified, or deleted during the migration process, the source server cannot guarantee migration integrity. |
| DMS | There is any activity on the source database during migration, and migration depends on the task scheduled during migration:<br>• A full load of existing data.<br>• The application contains cached changes.<br>• There is ongoing replication. | The migration must include database objects other than tables, primary keys and some unique indexes.<br><br>See *Limits for AWS Database Migration Service* in the AWS Database Migration Service *User Guide* at https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Limits.html for resource limitations of the DMS service. |

Use the `aws cp` command to move the files to an area that is local to your instance. You can migrate data to either S3 or EFS.

**Note**

EFS requires that your instance and the EFS system be in the same region. See https://docs.aws.amazon.com/general/latest/gr/rande.html#elasticfilesystem-region for information about available AWS EFS regions.

# 2 SAP Managed Data Services

If you need to provision a new SAP ASE database environment in the AWS cloud, or move your existing on-premises SAP ASE database environment to the AWS cloud, SAP Managed Data Services (MDS) has the practical knowledge and learned experiences that can help you successfully implement your SAP ASE solution in the AWS cloud.

The AWS Platform Transition service offered by SAP MDS will help you set up your AWS cloud platform, move your systems into the cloud, provide continuous improvements, and manage your operations through a well-established and tested Continuum of Care model. MDS will also help ensure your SAP ASE database meets SAP best practices and SAP requirements for the AWS cloud platform.

# 3 bcp Migration Method

Copy data out of the on-premises data center and back into the SAP ASE on AWS using the `bcp` utility.

Use the `bcp` migration method when:

- You are moving tables between the SAP ASE servers, or between an SAP ASE server and other data sources.
- The source and target database have different character sets.
- The source and target databases are on different platforms.

> **i Note**
>
> The source and target database versions need not be the same to use `bcp`.

Do not use this method when:

- You need to create the schema or metadata along with the data transfer.
- The data transfer must include database objects.

`bcp` does not create the database into which you are copying the data, nor does it copy out any objects (for example, stored procedures, indexes, and so on). Use the `ddlgen` utility first to create the database and the accompanying objects.

You must move the files that `bcp` creates from the machine running the source database to the machine running the target database.

Follow the steps in Using bcp to Migrate Data to AWS [page 6] to migrate databases from an on-premises database to a database running on AWS.

## 3.1 Using bcp to Migrate Data to AWS

The `bcp` utility copies data from the source table to a file.

### Context

The steps in this topic migrate the `pubs2` database.

## Procedure

1. Use the `ddlgen` utility to generate the SQL code to create the database and its objects. For example, the following command generates the SQL for the `pubs2` database on a machine named "big_machine" on port number 11724:

   ```
   $SYBASE/$SYBASE_ASE/bin/ddlgen -Usa -P -S big_machine:11724 -TDB -Npubs2 >
   pubs2_out
   ```

   > **i** Note
   >
   > `ddlgen` generates SQL code that creates a database the same size as the source database. If your target database uses a larger page size, you may need to edit the `ddlgen` output to increase the size of the target database. For example, the previous `ddlgen` output creates a 4 MB `pubs2` database (based on the source database size):
   >
   > ```
   > CREATE DATABASE pubs2
   >            ON master = '4M' -- 2048 pages
   > ```
   >
   > Edit the SQL code output to increase the size of your database, if necessary.

2. Use `bcp` to copy the data from the source database. Using the appropriate parameters, issue `bcp` once for each table in the database. `bcp` includes a lot of potential parameters. The parameters you include depend on the data you are copying out and the environment you are copying into.

   The examples below copy all the tables from the `pubs2` database from a server named "huge_server," and include these parameters:

   - `-U` – specifies the user name logging in to the server.
   - `-P` – password for the user.
   - `-S` – name of the server.
   - `-J` – requests the server to convert to and from `<client_charset>`. In this case, `iso_1`.
   - `-I` – location of the interfaces file
   - `-Y` – disables character-set conversion in the server. Instead it is performed by `bcp` on the client side when using `bcp out`.
   - `-c` – performs the copy operation using the `char` datatype as the default.

   ```
   ./bcp pubs2..au_pix out pubs/au_pix.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..authors out pubs/authors.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..blurbs out pubs/blurbs.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..discounts out pubs/discounts.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..publishers out pubs/publishers.bcp -Usa -P -Shuge_server -Jiso_1
   -I$SYBASE/interfaces -Y -c
   ./bcp pubs2..roysched out pubs/roysched.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..sales out pubs/sales.bcp -Usa -P -Shuge_server -Jiso_1 -I$SYBASE/
   interfaces -Y -c
   ./bcp pubs2..salesdetail out pubs/salesdetail.bcp -Usa -P -Shuge_server -
   Jiso_1 -I$SYBASE/interfaces -Y -c
   ./bcp pubs2..stores out pubs/stores.bcp -Usa -P -Shuge_server -Jiso_1 -I
   $SYBASE/interfaces -Y -c
   ./bcp pubs2..titleauthor out pubs/titleauthor.bcp -Usa -P -Shuge_server -
   Jiso_1 -I$bASE/interfaces -Y -c
   ```

```
./bcp pubs2..titles out pubs/titles.bcp -Usa -P -Shuge_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
```

3. Transfer the `bcp` files from the on-premises server to your SAP ASE instance running on AWS. Transferring the files typically uses the AWS S3 or EFS file storage using the following process:

| Storage Type | Description |
| --- | --- |
| S3 | To use S3 to copy to AWS storage, you need a correct access key and secret key. See the AWS documentation for details:<br>1. Upload your files to a bucket using the AWS S3 upload service located at https://s3.console.aws.amazon.com/s3/<br>2. Copy the files from your S3 storage to your EC2 instance using the AWS `aws cp s3` command. For example, the following command moves a zip file from the `pubs2` database `bcp` output to the release directory on EC2:<br>`aws s3 cp s3://bigbucket/ bcp_zip /work/AWS/` |
| EFS | Use *AWS Direct Connect* to mount your on-premises file systems to your EC2 instances. See *Amazon Elastic File System* at: https://aws.amazon.com/efs/ .<br>Your instance and your EFS must be in the same region. See http://docs.aws.amazon.com/efs/latest/ug/whatisefs.html . |

4. Create the database and its objects on AWS using `isql` to run the script created in the first step. For example, the following creates the `pubs2` database using the `pubs2_out` file after it has been transferred to a directory named `pubs` in $SYBASE:

```
$SYBASE/$SYBASE_OCS/bin/isql -Usa -P -S < $SYBASE/pubs/pubs2_out
```

5. Enable the `select into/bulkcopy/pllsort` database option in the database into which it is copying the data. This example enables this option in the `pubs2` database:

```
sp_dboption pubs2, "select into/bulkcopy/pllsort", true
```

6. Run a checkpoint in this database for the change to take effect:

```
use pubs2
go
checkpoint
go
```

7. Use `bcp` to copy in the data. This example copies in all the tables you copied out earlier to the `pubs2` database into a server named `big_cloud_server`:

```
./bcp pubs2..au_pix in pubs/au_pix.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
```

```
./bcp pubs2..authors in pubs/authors.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
./bcp pubs2..blurbs in pubs/blurbs.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
./bcp pubs2..discounts in pubs/discounts.bcp -Usa -P -Sbig_cloud_server -
Jiso_1 -I$SYBASE/interfaces -Y -c
./bcp pubs2..publishers in pubs/publishers.bcp -Usa -P -Sbig_cloud_server -
Jiso_1 -I$SYBASE/interfaces -Y -c
./bcp pubs2..roysched in pubs/roysched.bcp -Usa -P -Sbig_cloud_server -Jiso_1
-I$SYBASE/interfaces -Y -c
./bcp pubs2..sales in pubs/sales.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
./bcp pubs2..salesdetail in pubs/salesdetail.bcp -Usa -P -Sbig_cloud_server -
Jiso_1 -I$SYBASE/interfaces -Y -c
./bcp pubs2..stores in pubs/stores.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
./bcp pubs2..titleauthor in pubs/titleauthor.bcp -Usa -P -Sbig_cloud_server -
Jiso_1 -I$SYBASE/interfaces -Y -c
./bcp pubs2..titles in pubs/titles.bcp -Usa -P -Sbig_cloud_server -Jiso_1 -I
$SYBASE/interfaces -Y -c
```

# 4 Dump and Load Migration Methods

Using the `dump` and `load` commands to make copies of an entire database allows you to migrate data to AWS as well as between EC2 instances.

The `dump` and `load` commands copy the entire database, along with any database objects (for example, system procedures). The target database must already exist when you run the `load` command. The source and the target databases must both use the same page size.

Use `dump` and `load` commands when the device allocations on the target SAP ASE have equal or greater size than those of the source SAP ASE, and when the source and target databases have the same page size.

Do not use this method with the migration must load transaction dumps on different platforms, or when the dump from the source server uses a sort order that differs from the target server.

You must configure both instances for the `aws` command-line interface (CLI) to perform these tasks. See here for more information.

## 4.1 Using dump and load to Migrate Data to AWS

Make copies of an entire database by using the `dump` and `load` commands, which use the Backup Server to make the copies.

### Prerequisites

- The source and target databases use the same page size.
- The target database exists before you run the `load` command.
- Have a correct access key and secret key to perform the AWS `copy` command. See the AWS documentation at https://docs.aws.amazon.com/cli/latest/userguide for more information.
- Your instance and your Amazon Elastic File System (EFS) must be in the same region. See http://docs.aws.amazon.com/efs/ .

### Procedure

1. Verify that Backup Server is running on the server from which you are making the dump. If it is not running, use the `backupserver` binary (located in `$SYBASE/$SYBASE_ASE/bin`) to start it. For example:

```
./backupserver -SSYB_BACKUP -I/work/SP03_PL01/interfaces -e/work/SP03_PL01/
ASE-16_0/install/backup_error -M/work/SP03_PL01 &
```

See the *Utility Guide* for more information about starting Backup Server.

2. Perform a dump of the database. This example performs a dump of the `pubs2` database, making a copy in the `/work/Devices` directory:

```
dump database pubs2 to '/work/Devices/pubs2_dmp'
```

3. Transfer the database dump from the on-premises server to your SAP ASE instance on AWS. Transferring the files typically uses the AWS S3 or EFS file storage, using the following method:

| Storage Type | Description |
| --- | --- |
| S3 | To use S3 to copy to AWS storage, you need a correct access key and secret key. See the AWS documentation for details:<br>1. Upload your files to a bucket using the AWS S3 upload service located at https://s3.console.aws.amazon.com/s3/<br>2. Copy the files from your S3 storage to your EC2 instance using the AWS `aws cp s3` command. For example, the following command moves a zip file from the `pubs2` database `bcp` output to the release directory on EC2:<br>`aws s3 cp s3://bigbucket/ bcp_zip /work/AWS/` |
| EFS | Use *AWS Direct Connect* to mount your on-premises file systems to your EC2 instances. See *Amazon Elastic File System* at: https://aws.amazon.com/efs/ .<br>Your instance and your EFS must be in the same region. See http://docs.aws.amazon.com/efs/latest/ug/whatisefs.html . |

4. Log in to your SAP ASE on AWS using `isql`.
5. Load the database copy into SAP ASE. The following example loads the dump of the `pubs2` database:

```
load database pubs2 from '/work/AWS/pubs2_dmp'
```

## 4.2 Using dump and load to Migrate Data Between EC2 Instances

Make copies of an entire database by using the `dump` and `load` commands and use the `aws` command to copy the source database dump to an intermediary `S3` bucket, then use it again to copy this dump to the target EC2 instance.

### Prerequisites

- The source and target databases use the same page size.
- The target database exists before you run the `load` command.
- Configure of both instances for the `aws` command line interface (CLI). See https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-welcome.html .
- Have a correct access key and secret key to perform the AWS `copy` command. See the AWS documentation at https://docs.aws.amazon.com/cli/latest/userguide for more information.
- Your instance and your Amazon Elastic File System (EFS) must be in the same region. See the *AWS Command Line Interface User Guide* at http://docs.aws.amazon.com/efs/ .

### Procedure

1. Verify that Backup Server is running on the server from which you are making the dump (run `ps -ef | grep backupserver` to see the process). If it is not running, use the `backupserver` binary (located in `$SYBASE/$SYBASE_ASE/bin`) to start it. For example:

   ```
   ./backupserver -SSYB_BACKUP -I/opt/sap/interfaces -e/opt/sap/ASE-16_0/install/
   backup_error -M/opt/sap/ &
   ```

   See the *Utility Guide* for more information about starting Backup Server.
2. Perform a dump of the database. This example performs a dump of the `pubs2` database, making a copy in the `/opt/sap` directory:

   ```
   dump database pubs2 to '/opt/sap/pubs2_dmp'
   ```

3. Copy the files from your EC2 source instance to your S3 storage using the AWS `aws cp s3` command. For example, the following command copies the `pubs2_dmp` file to a bucket named `bigbucket` on S3 (this command is run from the directory that holds the dump file):

   ```
   aws s3 cp pubs2_dmp s3://bigbucket
   ```

4. Copy the files to your EC2 target instance from your S3 storage using the AWS `aws cp s3` command. For example, the following command, which is run from the directory that holds the dump file, copies the `pubs2_dmp` file from the `bigbucket` bucket on `s3` to a directory on the target EC2 instance:

   ```
   aws s3 cp s3://bigbucket/pubs2.dmp .
   ```

5. Log in to your SAP ASE on AWS using `isql`.
6. If necessary, create the database; `load database` does not create the database into which you are loading the dump. This command creates a 50 MB `pubs2` database on the default device:

   ```
   create database pubs2 on default = '50'
   ```

7. Load the database dump into SAP ASE. This example loads the dump of the `pubs2` database:

   ```
   load database pubs2 from '/opt/sap/data/pubs2_dmp'
   ```

8. Bring the newly loaded database online. For example:

   ```
   online database pubs2
   ```

   You are now ready to use the database.

# 5 sybmigrate Migration Method

Use sybmigrate when you need to migrate the database objects along with the data.

Do not use sybmigrate when there is any activity on the source or target server during the migration. If the objects are created, modified, or deleted during the migration process, the source server cannot guarantee migration integrity.

> **i** Note
>
> The source and target database platforms, database versions, and page size need not be the same to use `sybmigrate`.

Before you begin the migration process:

- Verify that your source server has a local server name by selecting from `sysservers`. The local server name should be defined in the `srvname` column. If it is not, issue the following `sp_addserver` command to set it as the local server, then restart the server for the change to take effect:

  ```
  sp_addserver <server_name>, local
  ```

- Define the entries for each server in the local and target server's interfaces file, so that the source server's interfaces file has an entry for the target server, and the target server's interfaces file has an entry for the source server.
- Define an empty version of the database you are migrating on the target server. For example, if you are migrating the `pubs2` database, there must be an empty version of a database named `pubs2` on the target server of the appropriate size. If you are migrating to a server that uses a larger page size, make sure the target version of the database is large enough to accommodate this larger page size.

## 5.1 Using sybmigrate to Migrate Data to AWS

`sybmigrate` allows you to migrate all data from the source database to the target, including database objects.

### Procedure

1. Issue the following command from the `$SYBASE/$SYBASE_ASE/bin` directory:

   ```
   ./sybmigrate -m setup
   ```

2. Select the source and target servers.
3. Specify the logins and passwords.
4. Click *Connect*.

5. Follow the on-screen instructions in the `sybmigrate` wizard.

> **i** Note
>
> If the migration fails, check the *Messages* tab for information. Re-run the migration once you address the cause of the failure. You need not restart `sybmigrate` to restart the migration.

# 6 DMS Migration Method

AWS provides the DMS utility for migrating data from your on-premises database to your AWS EC2 instance, or vice versa.

The DMS utility does not migrate database objects (for example, user-created stored procedures). How you migrate these objects depends on the engine type:

- If the source and target databases have the same engine types, use native utilities such as the SAP ASE `ddlgen` utility to export and import these objects to and from the target database.
- If the source and target databases have different engine types, use the AWS Schema Conversion Tool (AWS SCT) to generate these objects on the target database.

Use this method is there is any activity on the source database during migration, and migration depends on the task scheduled during migration:

- A full load of existing data.
- The application contains cached changes.
- There is ongoing replication.

Do not use this method if the migration must include database objects other than primary keys and some unique indexes.

See the *AWS Database Migration Service User Guide* for more information about source and target table limitations and prerequisites for migration:

- http://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SAP.html
- http://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.SAP.html

AWS provides complete documentation for the DMS utility. The general steps are described in Using DMS to Migrate Data to AWS [page 16].

## 6.1 Using DMS to Migrate Data to AWS

Use the AWS Data Migration Service (DMS) to migrate data securely from your on-premises servers to your EC2 instances.

### Procedure

1. Go to https://aws.amazon.com/dms/ .
2. Click *Get started with AWS Database Migration Service*.
3. If necessary, sign in to your Amazon account.
4. On the *AWS Management Console* screen, select *Migration and Transfer* > *Database Migration Service*.

5. Select the *Get Started* tab and click *Next*.
6. Enter your selections to create a replication instance and click *Next*:

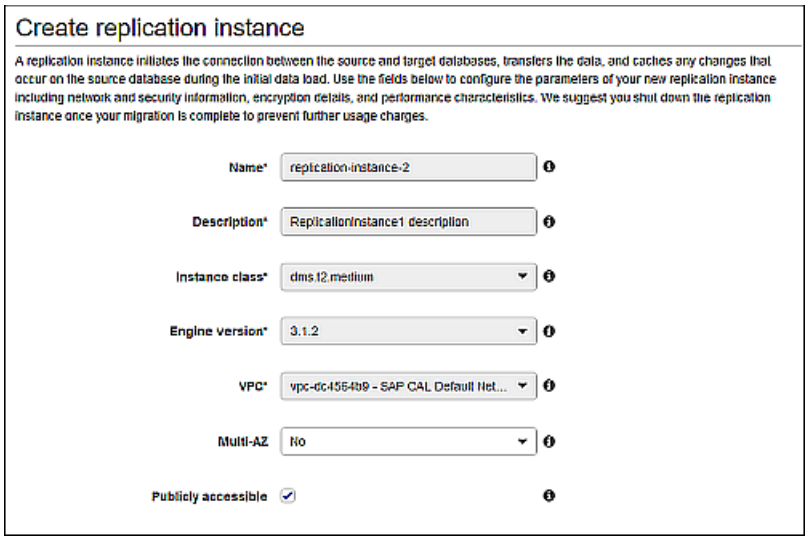


7. Enter your selections to create the source and target endpoints and click *Next*:

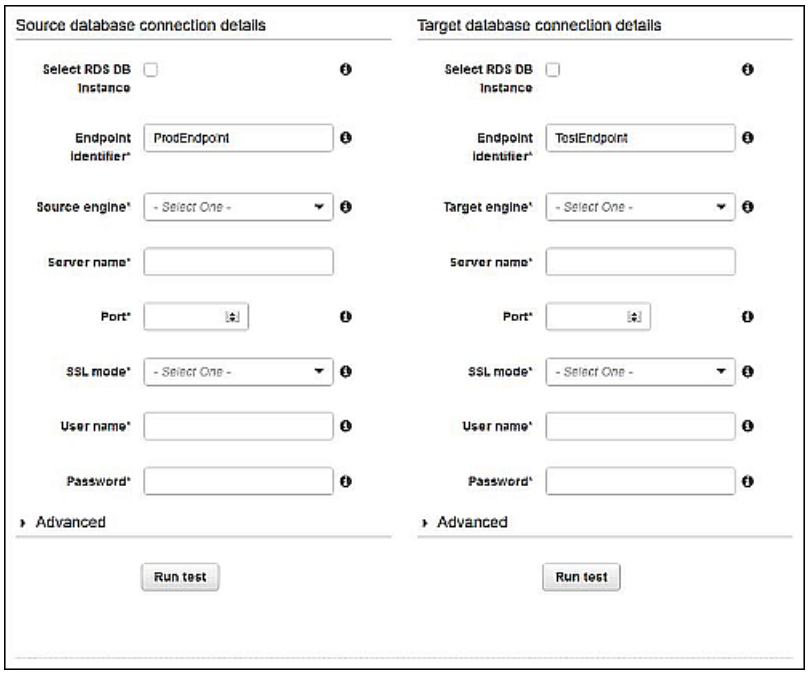


8. (Optional) test the connection by selecting your desired VPC and selecting *Run test*.
9. Enter your selectinos to create the task and click *Add selection rules*:

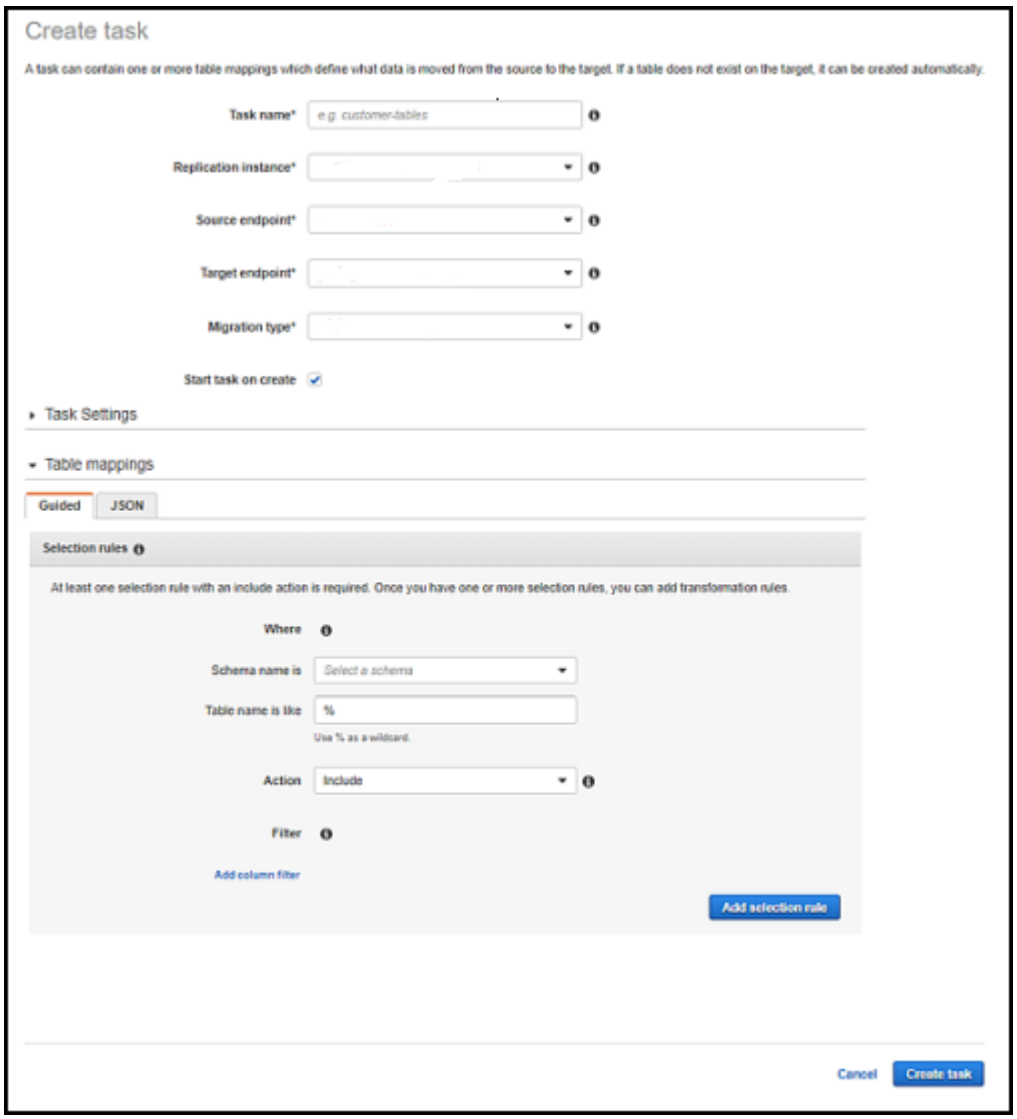


10. After creating the task, select the *Table Statistics* tab to see which tables that are loaded, as well as their details like schema name, load state, full load rows, and so on.

# Important Disclaimers and Legal Information

## Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

## Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

## Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

## Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

## Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap



THE BEST RUN SAP